

IBM Information

>>> On Demand

2007



Data Warehouse on system z – a Gold Consultant Briefing

Beth Hamel, hameleb@us.ibm.com &

Michaelyn Shelley-David, michaelyn@us.ibm.com



Act.Right.Now.

IBM INFORMATION ON DEMAND 2007

October 14 - 19, 2007

Mandalay Bay

Las Vegas, Nevada

Disclaimer

The information contained in this presentation has not been submitted to any formal IBM review and is distributed on an "As Is" basis without any warranty either expressed or implied. The use of this information is a customer responsibility.

The materials in this presentation are also subject to

- enhancements at some future date,
- a new release of DB2, or
- a Programming Temporary Fix (PTF)

IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, zSeries, System z.

THE FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC and WINDOWS 95.

For additional information visit the URL

<http://www.ibm.com/legal/copytrade.phtml> for "Copyright and trademark information"



Agenda

- Why System z? Why now?
- Operational Business Intelligence and System z
- Your feedback....guided Q and A





Why z/OS? *Why* now?

Act.Right.Now.



Value proposition of DW on System z

- **Qualities of Service**
 - Availability
 - Security and Regulatory Compliance
 - Scalability
 - Backup and recovery
- **Positioned for the future**
 - Web-based applications
 - XML support
 - Service Oriented Architecture (SOA)
- **Operational data and the ODS together means**
 - Reduced complexity
 - Reduced cost
 - Shared processes, tools, procedures
 - Streamlined compliance and security
- **Specialty engines improve TCO**
- **Better leverage System z skills and investment**



The Changing Warehouse Terrain

*“As a direct effect of the mixed workload, with continuous loading and the increase in automated transactions from the functional analytics in OLTP, the transactional DBMSs have an edge that challenges the DW DBMSs”
...from Gartner Data Warehouse Magic Quadrant*



Benefits of a transactional data server foundation

Optimized for real-time access,
High availability and reliability
Scalable, secure and auditable

Dedicated warehousing

Advanced data partitioning
Workload management



Cost of Ownership is King

The 'Hidden' Operational Costs of Computing

- Management and administration
 - ***'However, the costs of supporting and managing these complex environments and infrastructures have soared, and now far outweigh the customer's expenditure on new systems themselves'***
 - © Software Strategies 2005 11
- Security breaches
 - ***More Than 90% Of Companies Expose Sensitive Data***
Reconnex Insider Threat Index August 2005
 - ***Businesses Reluctant To Report Cyber Attacks***
2005 CSI/FBI Computer Crime and Security Survey
 - ***One In Four Identity-Theft Victims Never Fully Recover***
Nationwide Mutual Insurance Co. Survey July 2005
 - ***Card Associations Unite Setting Standards to Fight Fraud***
Green Sheet Inc. August 2005 Issue 2
- Downtime
 - ***Cost of downtime can vary by industry and can range from hundreds of thousands to millions of dollars per hour***
 - ©Robert Francis Group. All Rights Reserved 2005



IBM Mainframe solutions are highly available, highly secure and highly managed to help lower TCO



Security - Empowered by System z

- ✓ Built in security to address the full spectrum of security requirements
- ✓ Policy based security management
- ✓ Addresses security functional domains
- ✓ Designed to meet evolving applications needs
- ✓ Addresses multiple resource types- applications, data, networks
- ✓ Positions System z as a “secured vault”
- ✓ Proven heritage and experience base
- ✓ Common criteria certified
- ✓ Simplifies security infrastructure
- ✓ Integrated System z, z/OS and DB2 security
- ✓ Most important -----
 - ✓ DB2 is integrated with it!
 - ✓ It Works!



“Whilst the performance and resilience characteristics are formidable, it is the security features that are likely to attract most attention”


Tony Lock – Chief Analyst, Bloor Research 2005

**Proven secure by 40
years of operation!**



System z makes data available

- **System z and DB2 for z/OS provide highest levels of availability, resiliency, security, and recovery**
 - **Sysplex and DB2 data sharing**
 - Availability (GDPS for DR)
 - Parallelism and Horizontal scalability
- **System z Work Load Manager**
 - Can balance workload with different SLAs based on business policies/needs
 - Provides 100%+ system utilization
- **Single System**
 - Easy to integrate
 - Less costly to manage, Lower TCO
- **Provide true real-time Operational Data Store (ODS)**
 - Operational data is on DB2 for z/OS
 - Real time ODS can be on same DB2
 - Near real time ODS kept in sync with the operational data
 - (same or different DB2 subsystem)




Technology Evolution with Mainframe Specialty Engines

★ Building on a strong track record of technology innovation with specialty engines, IBM is introducing the System z9 Integrated Information Processor



IBM System z9 Integrated Information Processor (IBM zIIP) 2006

Centralized data sharing across mainframes



Internal Coupling Facility (ICF) 1997



Integrated Facility for Linux (IFL) 2001

- Support for new workloads and open standards



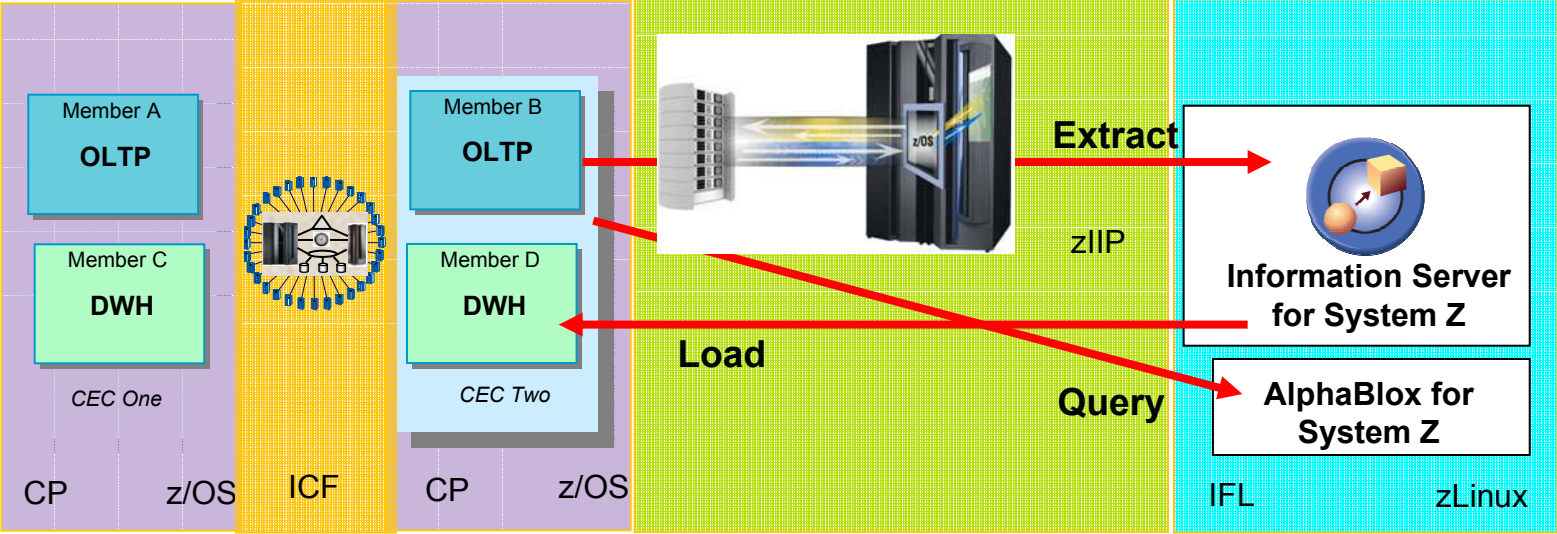
System z9 Application Assist Processor (zAAP) 2004

- Incorporation of JAVA into existing mainframe solutions

- Designed to help improve resource optimization for eligible data workloads within the enterprise



Specialty Processors in a DB2 for z/OS Warehouse Solution



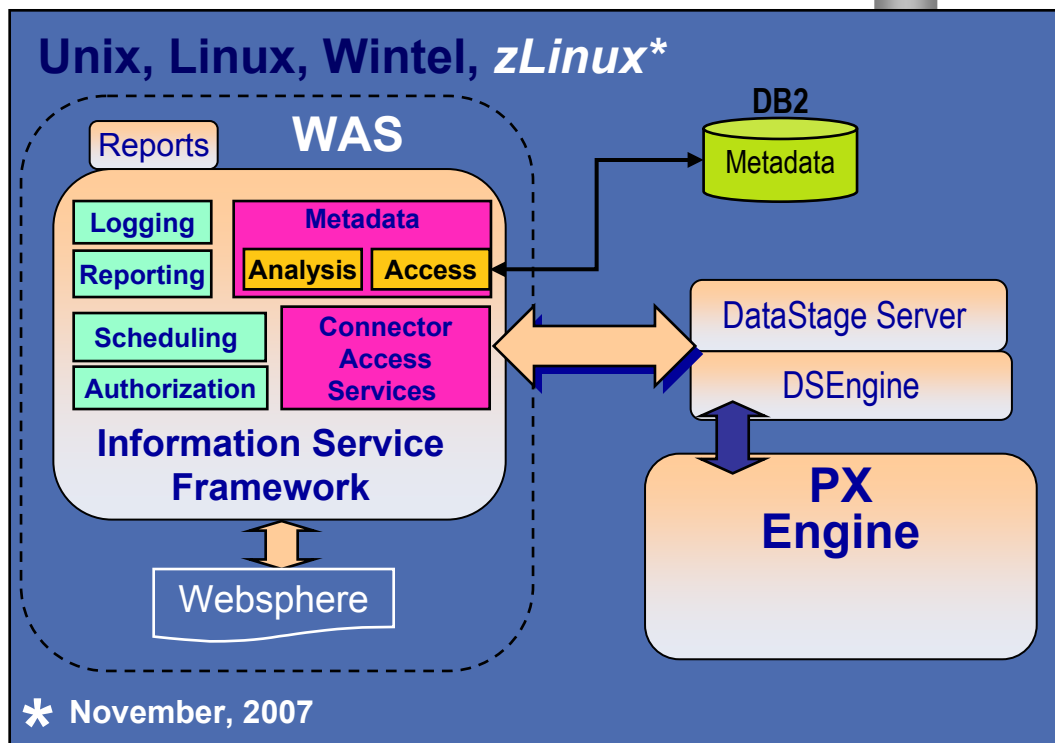
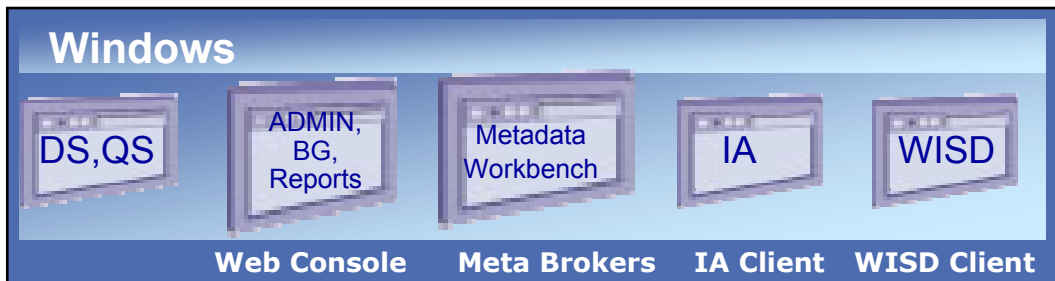
ICF – Uniquely allows a Data Warehouse database to coexist with an OLTP database

- IFL – Enables efficient data movement (secure, high-speed hipersockets)**
- Lowers TCO through reduced hardware and software costs
 - Enables use of zIIPs during extract and further reduces costs

zIIP – Further enables lower cost of Business Intelligence queries



IBM Information Server for System z



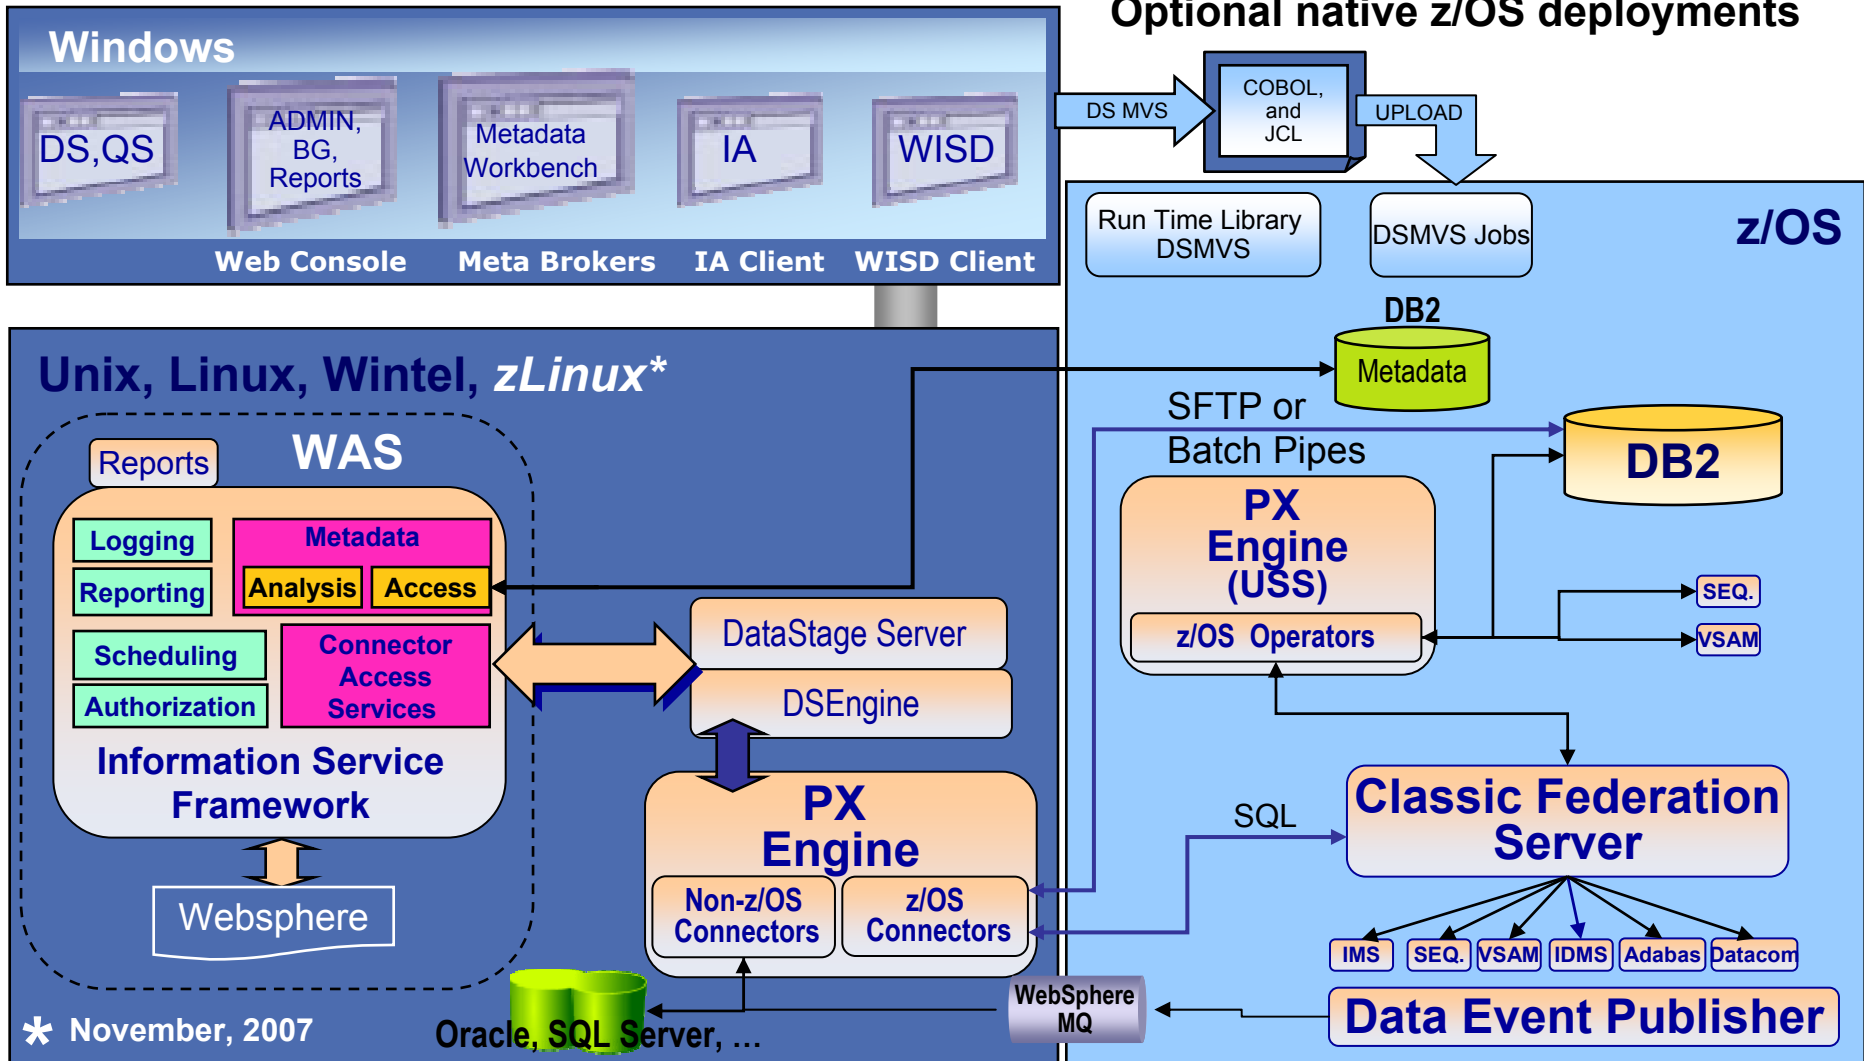
* November, 2007

New Linux for z deployment option

- Robust, parallel processing
- Hipersocket connectivity to z data
- Full Information Server suite: Information Analyzer, QualityStage, DataStage, Information Services Director, ...
- Minimal impact on z/OS costs: Leverages IFLs and zIIPs



IBM Information Server for System z



* November, 2007

Oracle, SQL Server, ...

WebSphere MQ



IBM Information Server for System z

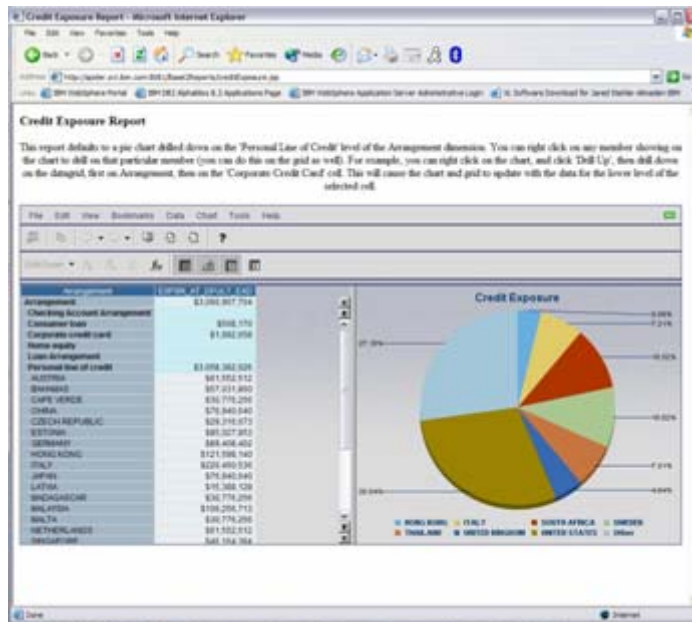
Benefits of this hybrid architecture

- Significant cost savings
 - z/OS MIPs consumption dramatically reduced vs. USS or MVS approaches
 - Minimizes impact on other z/OS software costs
 - All Job Processing is on zLinux (except the z/OS data access)
 - MIPs charged at IFL rate ... NOT z/OS rate
 - DB2 workload on z/OS can qualify for offload to ZIIP specialty engines
- High performance z data connectivity
 - Batch Pipes for DB2 load, DRDA to DB2 over hipersockets
 - SQL to Classic over hipersockets
 - Integration with MQ and therefore with the Data Event Publishers
- Seamless integration with other IBM Information Server platforms
 - Same operational architecture and metadata Repository
 - Eliminates deployment issues
 - Maintains value of DataStage for z/OS investments



Embedded Analytics with DB2 Alphablox

- Provide reporting and dashboard capabilities on key indicators
- Leverage operations such as ranking, ordering, filtering, trending, and other sophisticated statistical functions and calculations
- Drive data analysis from multiple data sources, both relational and multidimensional



Available since March 2007: IBM DataQuant

- Adds compelling new Warehouse/Business Intelligence component to WH on z
 - Visual Dashboards, Enhanced Graphical Reporting & Analytics, Security and Personalization, SOA Layer
 - Adds over 100 analytical functions to queries, reports and visual solutions
- Extends page-based QMF 'visual reports' with a comprehensive array of charts and graphical objects
- Transforms text-based QMF reports into highly graphical, information-rich documents



CLIENT PROFITABILITY AND ROI DASHBOARD

Dashboard | Rank Analysis

Account Selection
Region / WP
All Regions
All (All Regions)

Dashboard / ROI
All (All Regions)

ROI Analysis
All (All Regions)

Health Interval
All (All Regions)

ROI Analysis
All (All Regions)

ROI Analysis
All (All Regions)

Account	Oct 2005	Nov 2005	Dec 2005	Annual 2005	Oct 2006	Nov 2006	Dec 2006	Annual 2006
Accounting/Finance/HR/IT	26.03%	3.21%	25.01%	18.1%	29.54%	45.11%	34.14%	31.11%
Accounting/Finance/HR/IT	11.05%	5.36%	4.68%	5.01%	134.48%	175.71%	23.05%	111.08%
Accounting/Finance/HR/IT	17.05%	4.01%	21.58%	11.04%	37.04%	76.21%	23.15%	30.14%
Accounting/Finance/HR/IT	5.05%	11.06%	13.7%	11.04%	13.11%	110.21%	20.14%	110.14%
Accounting/Finance/HR/IT	10.05%	11.06%	11.06%	5.06%	10.14%	100.14%	11.04%	11.04%
Accounting/Finance/HR/IT	0.05%	0.14%	10.14%	10.14%	11.11%	100.14%	17.14%	10.14%



DB2 V8: More Than 50 Features Relevant to BI

Performance

- Data-partitioned secondary indexes (DPSI)
- Multiple DISTINCT clauses in SQL statements
- Reduced lock contention on volatile tables
- Coupling Facility lock propagation reduction
- Multi-row INSERT/FETCH
- REOPT(ONCE) to reduce host variables impact on access paths
- Index-only access for VARCHAR columns
- Backward index scan
- Faster short PREPARE
- IN access path performance
- DDF performance enhancements

Business warehouse

- Sparse index for star join
- More tables in join
- Common table expressions
- Recursive SQL
- Materialized query tables

Continuous availability

- Changing clustering index as online operation
- Elimination of BUILD2 phase of REORG with DPSIs
- Online schema evolution for many column types
- Volume-level, automated backup and recovery
- CI size larger than 4 KB
- More log data sets
- Conditional restart enhancements
- Support for synchronizing log point

Architecture

- Unicode support
- Introduction of DB2 Connect
- DB2 Universal Driver for JDBC
- 64-bit virtual storage for most DB2 storage areas
- Up to 4096 partitions
- Longer table/column names
- SQL statements up to 2 MB
- ASCII precompiler

Ease of use

- Clustering decoupled from partitioning
- New REORG option to reorganize all partitions in Reorg-pending state
- CREATE INDEX invalidates statements from dynamic statement cache
- Indexes created as deferred are ignored by DB2 optimizer
- LOB ROWID transparency
- Collecting distribution statistics on arbitrary sets of columns with RUNSTATS
- Fast cached SQL invalidation
- Automatic space management
- Statements IDs of cached statements as input to EXPLAIN
- Statement ID in IFCID 124
- Long-running non-committing reader alerts
- Lock escalation reporting
- Transaction-based DB2 accounting and workload management
- Stored procedures to facilitate database administration
- Network statistics with DB2 Connect
- DRDA ping
- Comments in dynamic SQL
- CTE-based optimizer hints



DB2 9: Another Feature Rich Release for BI

Performance

- New row internal structure for faster VARCHAR processing
- Fast delete of all the rows in a partition
- Numerous enhancements in 'smaller' LOB performance
- Fast LOB streaming
- Reducing log latch contention
- Deleting first n rows
- Skipping uncommitted inserted/updated qualifying rows
- Faster release of LOB locks
- Reducing data sharing overhead for global indexes
- Functional indexes

Business warehouse

- Dynamic index ANDing
- Reduce temporary tables materialization
- Generalizing sparse index/in-memory data caching

Continuous Availability

- Partition-by-growth as a means to remove non-partitioned tablespace size limit
- Full support for system-level backup and recover (automatic offload to tapes and individual objects recovery)
- Renaming SCHEMA and VCAT to facilitate fast database provisioning
- Rename index
- Reorganization of LOBs to reclaim space
- Online REORG enhancements
- Online REBUILD index

Architecture/SQL

- Thin DB2 Connect Client
- FOR BIT DATA collating sequence (VARBINARY)
- Full JDBC compliance
- Enable Decimal Float data type (pre-conditioning)
- BIGINT data type
- Index compression

Architecture/SQL (con't)

- Provide more VS relief for thread related storage (partially)
- Unicode support for all CLI functions
- MERGE statement
- SET operations

Ease of Use

- Implicit objects creation
- Enhancing real time statistics (Optimization Service Center)
- Autonomic reoptimization
- Integration of Real Time Statistics tables into the catalog
- Simulating indexes in EXPLAIN (Optimization Service Center)
- More autonomic bufferpools tuning (WLM synergy)
- RLF support for end-user correlation
- TRACE support for end-user correlation
- Enhance tracing in DB2 Connect
- Identifying unused indexes
- Enhancing IFC for IRLM diagnostics
- DSNACCOR enhancements



Another Motivation – Going Green





Operational BI and z/OS

Act.Right.Now.

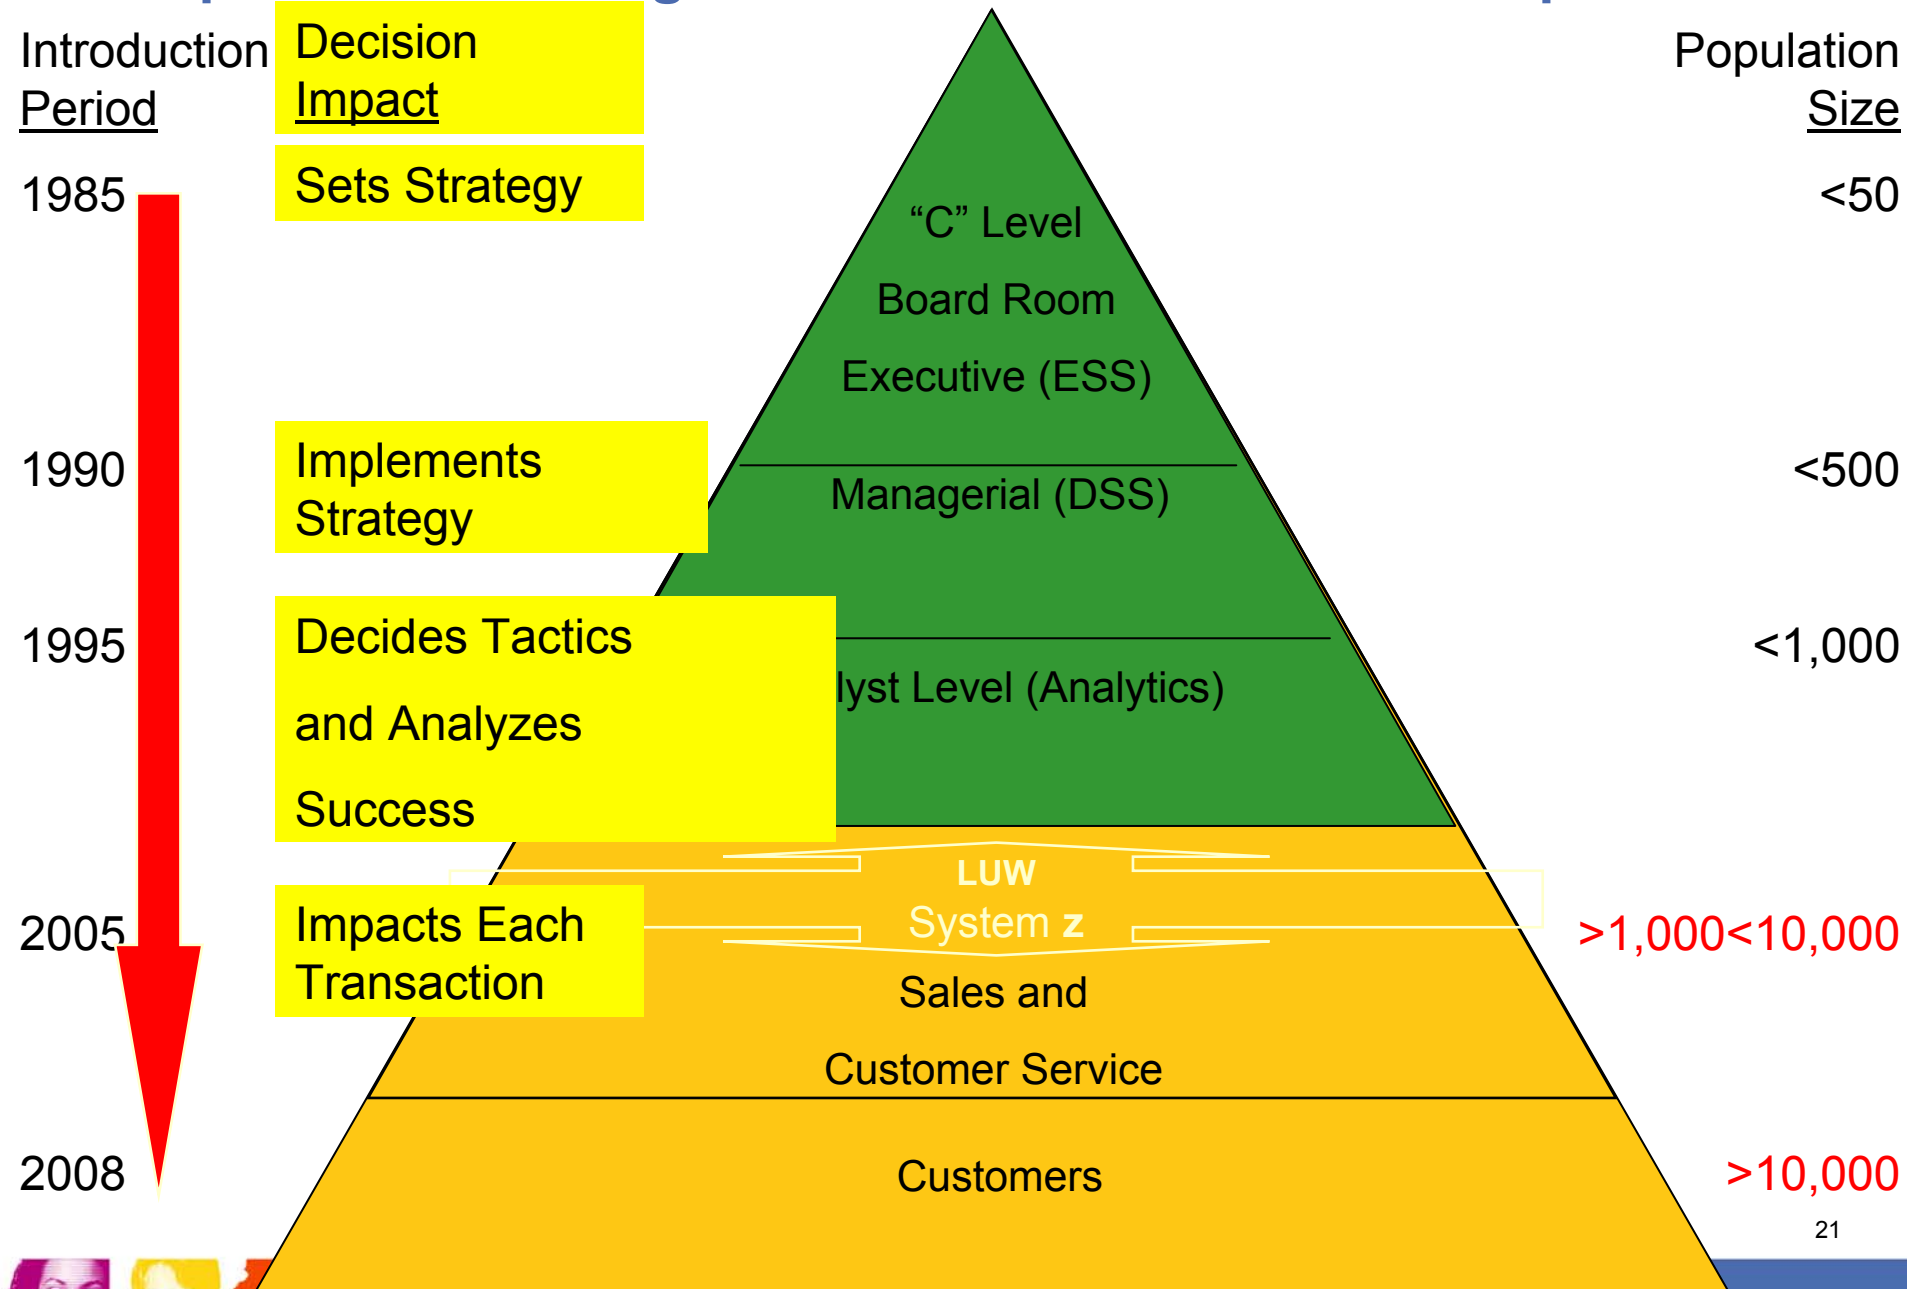


What is Operational Intelligence ?

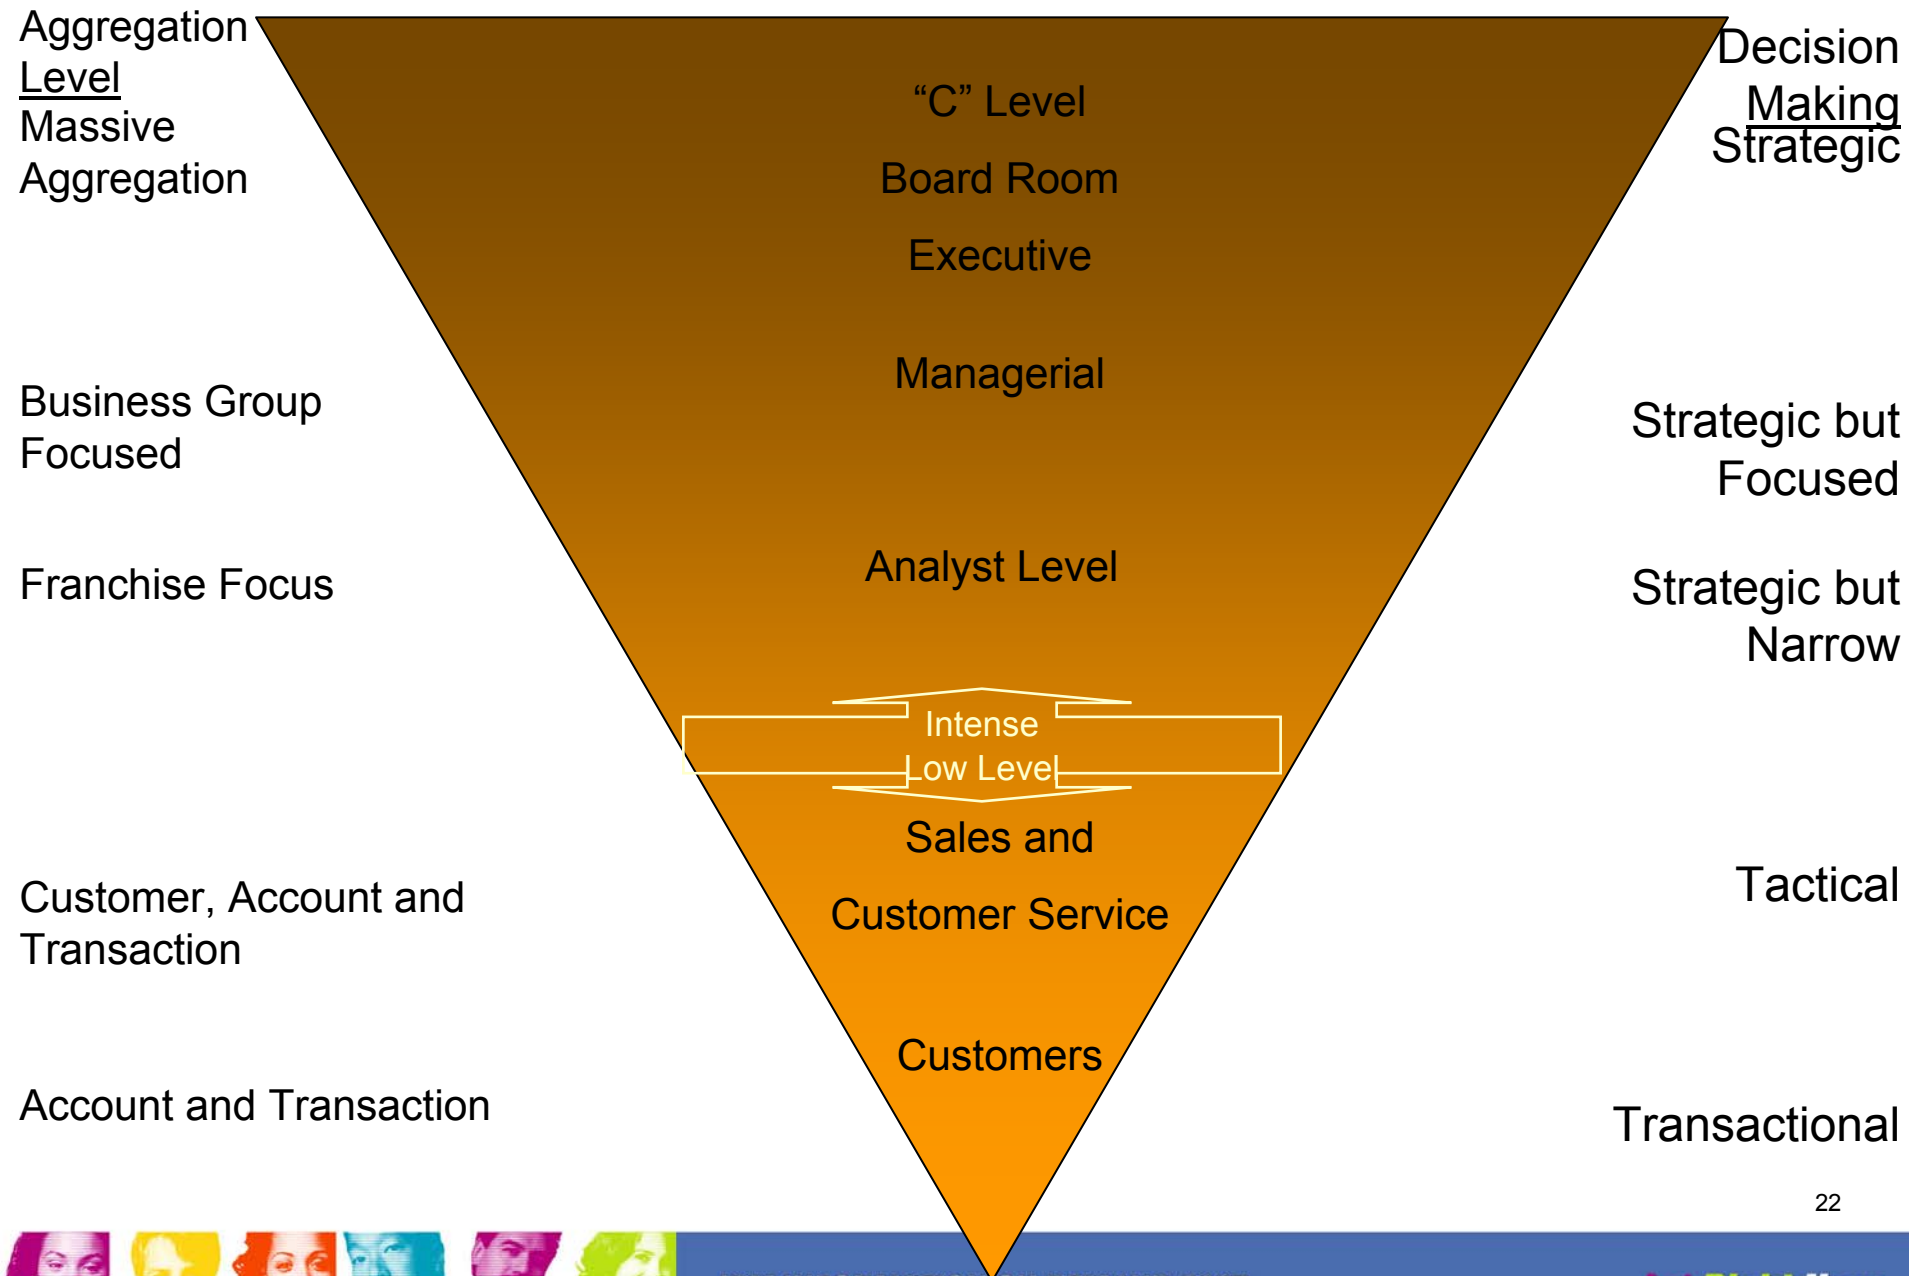
- Major Growth Area in Business Intelligence
 - Deliver BI content
 - to customer facing people to optimize
 - Sales
 - Customer Service
 - to improve Corporate Efficiency
 - and improve customer retention (i.e. “stickiness”)
 - via components within operational systems & information portals
 - “Embedded Analytics” provides guided analysis and efficiency
- Customer Service/ Self Service Examples
 - Wireless Customer Service
 - Credit Card Self Service Reporting
 - Consumer Products Packaged Goods Sales



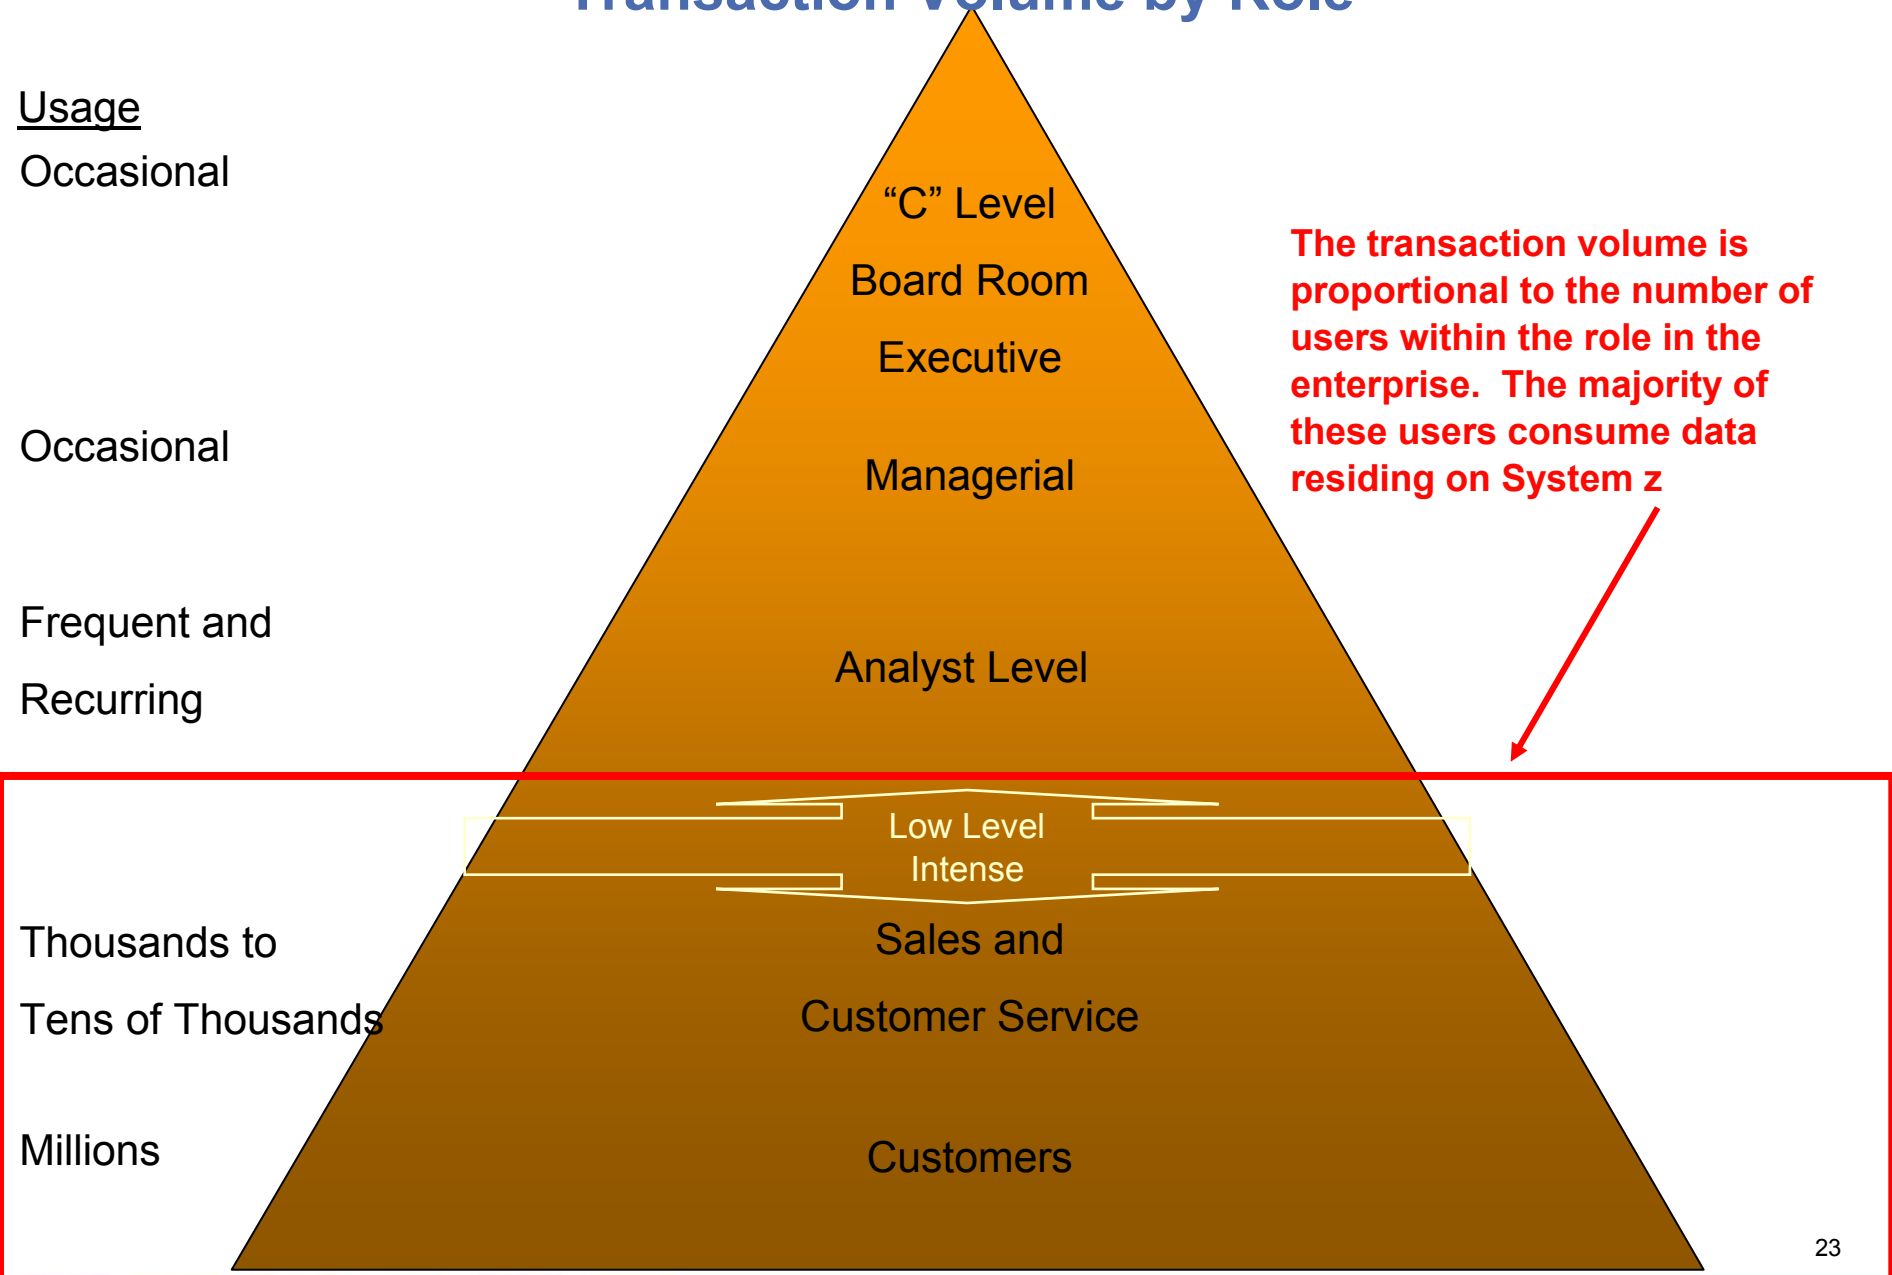
Operational Intelligence Introduction and User Population



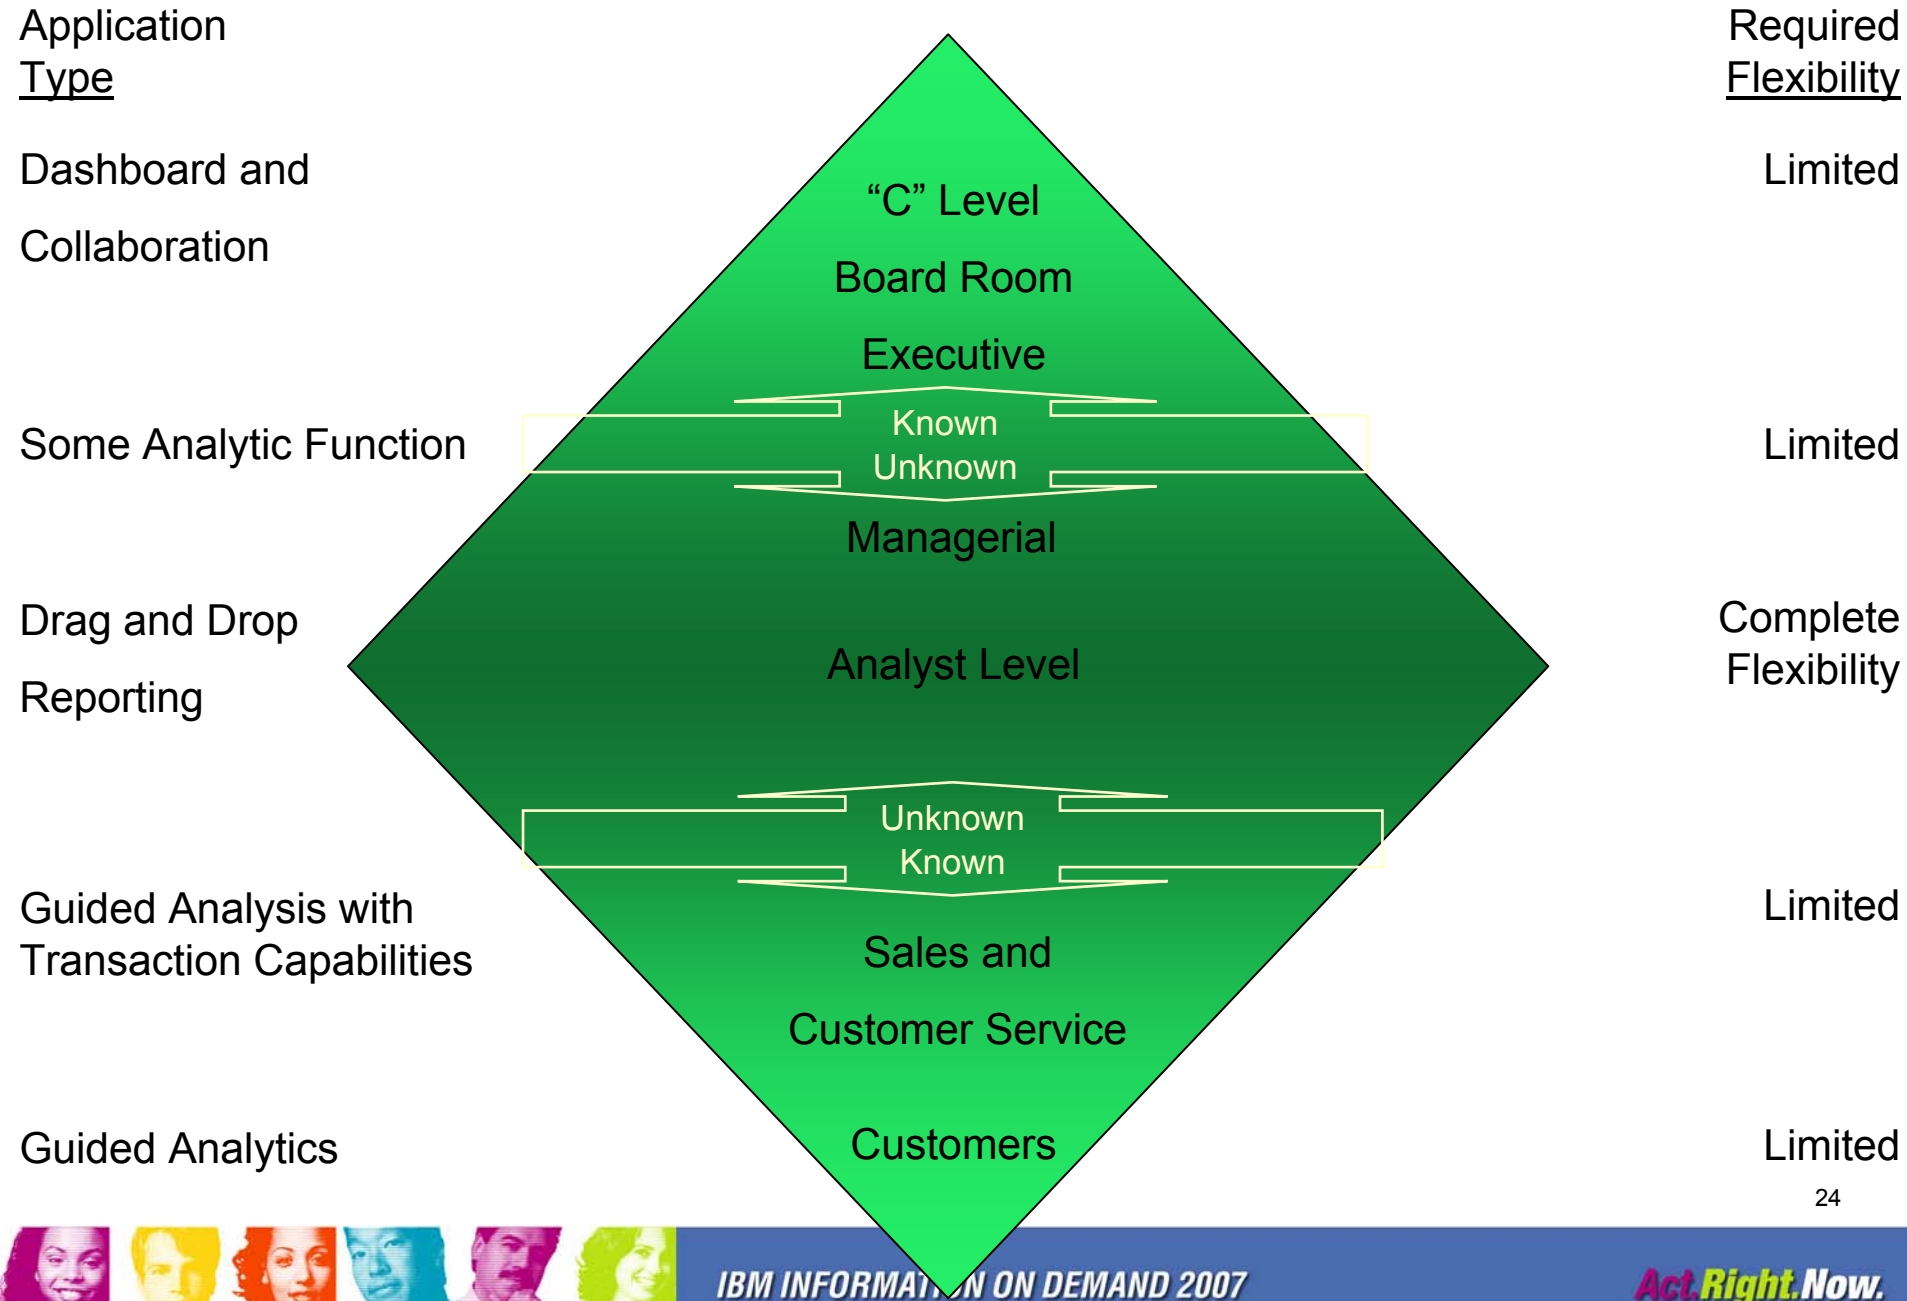
Aggregation Level and Decision Making



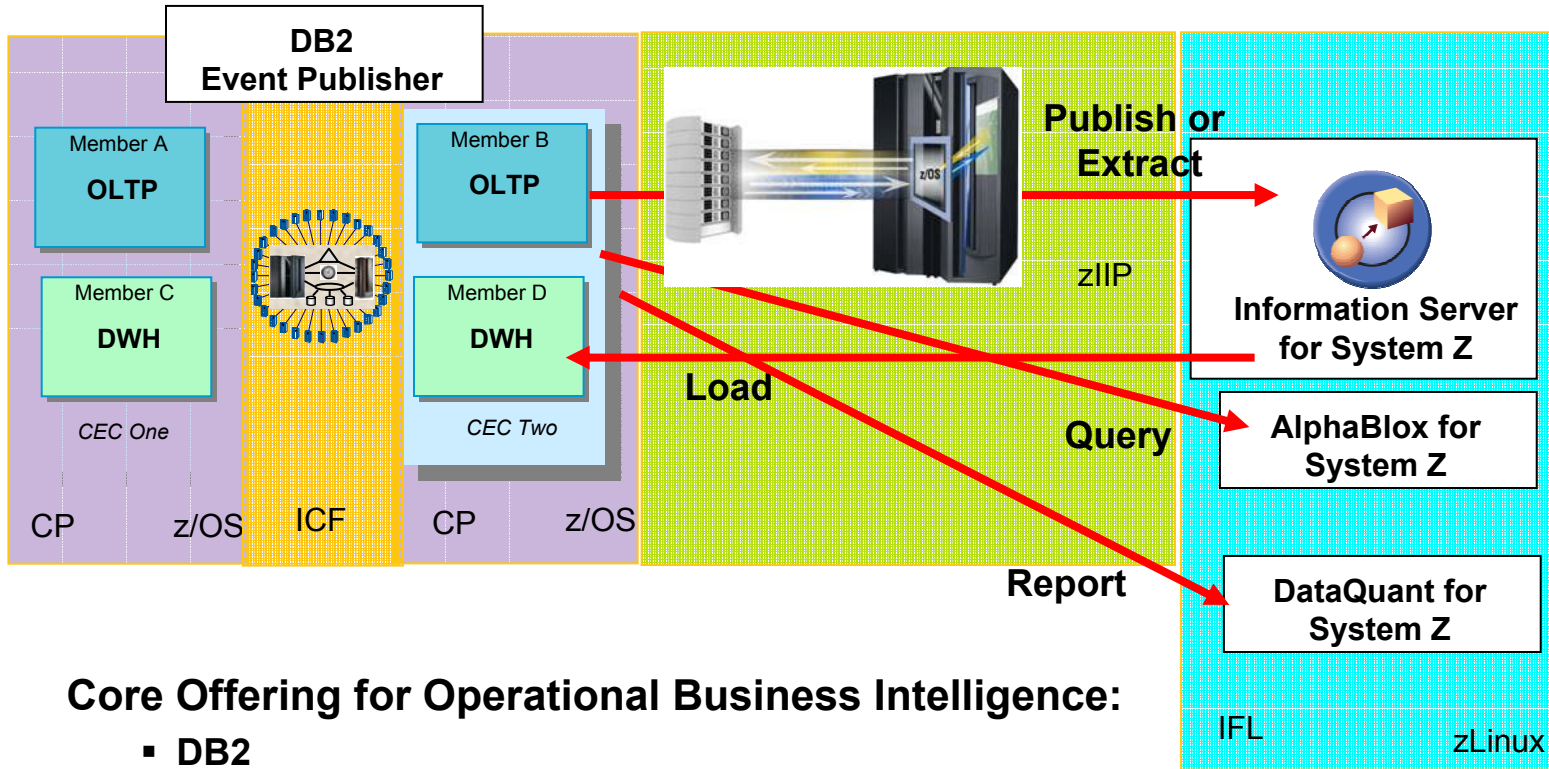
Transaction Volume by Role



Operational Intelligence – Application Flexibility and Workload



DB2 for z/OS Warehouse - OI Solution



Core Offering for Operational Business Intelligence:

- DB2
- DB2 Event Publisher
- DataStage on zLinux
- AlphaBlox on zLinux
- DataQuant on zLinux

